FEC Frothingham Electronics Corp.

Testing Thermal Response of Diodes

If there is a single most important thing that I can tell you about Thermal Response testing, it is this: know your product and set your own test limits accordingly!

Industry wide specifications are certainly useful and set a standard that no credible product should fail to meet. However it is entirely possible that a bad part by one manufacturer (process/design) may test better than a good part by another.

How could that be?

The thermal properties of a given diode design are limited by the geometry and metallurgy of the diode package and the semiconductor chip within it. Given a certain design, if all of the parts are assembled perfectly and perfectly bonded together, the part will have predictable thermal properties.

A part designed with a slightly different geometry or different metal or alloy in any of its parts would have different thermal properties if processed perfectly.

Both parts may be suitable for the application and be reliable if assembled perfectly (or nearly so). Suppose the measured response of a particular part from the "best" design varies significantly from its theoretically perfect reading. It is highly likely that there is a poor bond within that part.

Parts that have detectable imperfect bonds early in their life have a much higher probability of premature failure!

A part from a design with a slightly poorer initial thermal performance, but perfectly bonded (or nearly so) will be much more likely to survive the stress of repeated heating and cooling that diodes are subject to in typical applications.

What to do

Thermal Response testing with carefully selected test conditions and limits can detect defects quickly and non-destructively. Before getting into details of setting up a test program, I would like to share the following typical case history that I have experienced several times in my years of testing Thermal Response.

A manufacturer is experiencing excessive returns of rectifiers with catastrophic failures. The parts are cut open to reveal poor bonding of the chip or cell to the package. If no Thermal Response tester is available, the engineers try various tests such as high current SURGE pulses or tight VF limits. It doesn't work!!

We supply one of our Thermal Response testers and help our customer to evaluate the product. Often we pick out parts that have fairly small deviations from the product norm. These parts are opened up and, sure enough, bad bonds.

But this is only part one of the story. Now that the means of evaluating the product is available, the manufacturing process can be adjusted and the results monitored. The optimum process settings can be found and in the end the process norm is shifted significantly toward the best that the design will allow. Problem solved!

Finding the bad parts. A quick overview

(long-winded justification to follow)

Make sure all the samples are good parts electrically.

Electrical rejects may produce unreliable thermal readings which will mess up your statistics. Eliminate the electrical rejects FIRST.

Use DVF, not THETA

THETA reads in Degrees Celsius/Watt and DVF reads in Delta VF. DVF works better.

Find the optimum SAFE value for measurement delay (TMD)

Use the plot function "Cooling Plot" for every test program and for every part number and (if applicable) from every manufacturer. This allows you to select a low TMD but not so low as to be invalid.

Use a higher rather than lower IM

10mA is a good general purpose value although some large slow rectifiers may need more and some published specifications call for less. For internal statistical analysis, you can choose.

Collect data to plot histograms at several pulse widths (TH)

I would suggest 10mS, 50mS, and 250mS. Bad parts may show up best at one or more of these values.

Use a test current at each pulse width that produces a DVF of roughly 50mV to 100mV

For silicon junction rectifiers that would be a delta temperature of roughly 25 to 50 degrees Celsius. Schottky rectifiers would get a bit hotter but still ok.

They may ALL be bad (sort of)

If you are suspicious of this population already, consider the possibility that although some parts are worse the best ones could be better.

The long winded justifications (please read on)

ELECTRICAL TESTS FIRST

This is intended to apply to statistical analysis of Thermal Response, not overall acceptance testing. If you include Thermal Response in a test program including electrical tests, you could test Thermals first if you want. This would ensure that in the unlikely event that the part was damaged by the Thermal test, it would be rejected by the following electrical tests (especially the IR test).

Please keep in mind that the Thermal test leaves the junction temperature somewhat higher than the ambient. If the pulse width is on the high side, such as 250mS, it make take a few seconds to cool. The 10mS end of the range would of course cool much faster.

DVF vs THETA

If your final test is controlled by some published specification requiring reading Degrees/Watt then you should use THETA. You can still do your internal analysis as DVF and establish correlation between the DVF limits you establish and the THETA limits in the final program. You will find the correlation pretty good but not perfect and here is why.

At any fixed operating current, a rectifier will get hotter if the VF is higher. This will make the DVF reading higher also. The THETA test compensates for this difference by calculating the rise in temperature per Watt.

However a part operating at a higher temperature is more likely to fail! Also within any given lot of parts made by the same process, the part with the high VF is more likely to have a defect such as poor bonding or a cracked chip. Why compensate for that?

Actually if the part is a zener diode which would normally operate in the "Reverse" direction, there may be some justification for this compensation. You decide.

Find the best TMD

Everyone would like to use the lowest possible value for TMD. The reason for this is that the junction begins to cool immediately as the heating pulse current is turned off. The DVF measured will be slightly (to considerably) lower than the true value depending on the value of TMD used.

However, the lowest value that can be used is not constant. It is effected by several properties of test equipment which vary with the heating current (IH), the measurement of current (IM) and even the wiring to the test fixture and the fixture itself.

In addition, the lowest safe value is also effected by the carrier recombination time of the diode being tested. Fast Recovery diodes can usually use a lower value than the standard recovery types.

If the value used for TMD is marginally too low the results may appear reasonable but in fact be false and misleading.

Be safe. Always use the Cooling Plot function which will plot the Thermal test against TMD with the other test conditions programmed to the values you will be using. The readings plotted for moderate to long TMD values will be correct. As the TMD gets smaller the readings will go up in a smooth curve. At some point the curve may depart from this smooth predictable path. Use a TMD which is safely higher than this point.

Generally you can use a lower TMD if you use a higher IM (within reason). 10mS is a good general purpose IM for most small to medium power diodes. Large, slow rectifiers may benefit from a much higher IM (perhaps 100mA).

Try several pulse widths

I assume that you are trying to detect the bad bonds that can lead to premature diode failures. Diodes will have two (or more) internal bonds at various physical distances from the junction. The amount of time needed for heat generated at the junction to pass through a given bond is related to that distance.

There are many ways to mount a silicon chip within a diode package and so the distance to each bond may vary from part number to part number and from manufacturer to manufacturer.

You could use a very long pulse to be sure that the heat passed through all of the bonds, but that is not necessarily the best. If the problem bond is very close to the junction a long pulse will heat the diode structure well beyond the bond and give a higher reading. This higher reading not only dilutes the small variations due to the problem bond but may also be effected by other normal variations not related to the problem.

Collecting data for several pulse widths allows you to find the optimum width for the particular diode design being evaluated.

Selecting the right heating current (IH)

Since the variation in DVF that you are looking for may be fairly small, you would like to have a reading that is large compared to the normal repeatability errors in the equipment. A reading of 100mV is well out of the noise and is a good choice. Silicon junction rectifiers have a K factor of roughly 0.5 degrees per mV (we express it as 1/K which would be 2.0mV per degree).

The current (IH) needed to get a DVF or 100mV will be higher for low values of TH and vice versa.

A delta temperature of 50 degrees Celsius is a good safe value. That would give a final junction temperature of 75 degrees if the ambient is 25 degrees. You could go a bit higher if necessary, but avoid very high temperatures as they could distort the data or damage the diode.

Also keep in mind that Schottky diodes generally have lower K factors.

Analyzing your histograms

Let us assume that you are already suspicious of the quality of the sample being tested. I am assuming also that you have removed any electrical rejects.

If the sample contains good parts with only "normal" manufacturing variations you should see a "normal" bell curve with a Sigma of (dare I say?) well, not too many %. If you see "flyers" on the high side, these are probably parts with significantly bad bonds.

Flyers on the low side however show a bigger problem. For any given diode design, it is impossible for the process to be more than 100% perfect (of course). If there are flyers on the low side, these are probably parts that are at least approaching perfection and the main distribution is centered on something that well let’s say could use some improvement.

We at FEC have been through this process a number of times over the years and we stand ready to offer assistance if you wish to call on us.

Testing Thermal Response of Diodes, Part 2

THERMAL RESPONSE is a non-destructive electrical test for evaluating the quality of the bond between a diode junction and the package in which it is mounted. It is most commonly used to check the solder bond between the semiconductor "chip" and the package. Solder voids in this bond cause poor thermal conductivity leading to failure of the device due to over-heating of the junction. A moderate amount of voids has little or no effect on the electrical properties of the part, so attempts to detect them by changes in voltage drop generally fail.

The THERMAL RESPONSE test measures the increase in junction temperature due to a "heating pulse" of specified current and duration. The current and pulse width are chosen to heat the junction moderately and to allow time for heat to travel through the bond but not much further. A pulse that is too long will allow heat to travel into the leads or mounting surface. This results in a test that is sensitive to the thermal connection between diode and its test socket or heat sink. Conversely, a pulse that is too short will not allow heat to travel through the bond and so the result will not be effected by the quality of the bond.

Typical pulse widths are in the range of 10 to 50 milliseconds, although shorter or longer pulses may be used for some purposes. Once the pulse width is chosen you must choose an appropriate heating current (IH). The value is not critical but you must avoid two extremes.

At the low extreme the heating of the junction is so low that the DVF is very small and difficult to measure. At the high extreme the device is heated so much that it may be damaged or the K factor may change significantly making the results unreliable. A good choice would be a current that causes a rise in junction temperature of 40 to 50 degrees Celsius. This causes the DVF to be typically in the 70 to 100MV range (for a single Silicon Diffused junction) which is easy to measure and a junction temperature of 75 degrees will not damage the device or cause a significant change in K factor.

The usual first question asked by a new user is "how do you measure the junction temperature". The most common method is to make use of the known relationship between junction temperature and VF (Forward Voltage) at a low current.

The current is usually in the 1 to 10mA range and is referred to as "IM". This relationship is referred to as K or 1/K. K is the ratio of the change in temperature over the resulting change in VF. This might typically be 0.5 degrees/mV. Expressed as 1/K that would be 2.0mV/degree. The "K Factor" is usually reasonably constant over a lot of diodes of the same type.

It is common to measure K on a sample and then to use the average value for all of the parts in the lot. K is measured by reading VF at room temperature and at a moderately elevated temperature in an oven or hot bath. The test current (IM) should be the same as used for the THERMAL RESPONSE test.

K is not constant over a wide range of currents.

There are two variations on the THERMAL RESPONSE test. The tests are commonly referred to as DVF (Delta VF) and THETA. I will describe DVF first since it is the most fundamental of the two.

First a measurement is made of VF at IM and the reading is remembered.

Next a constant current pulse of the specified amplitude (IH) and duration (TH) is applied.

Finally, the current is set back to IM and after a brief delay (TMD) the VF is measured once more and then subtracted from the initial cold VF. This results in the test result which is DVF. It is not necessary to know K when doing a DVF test.

THETA is almost exactly the same electrical test except that the VF is also measured at the heating current IH. This allows us to calculate the heating power (VF x IH). We now have enough information to express the results in "Degrees per Watt". THETA in Degrees per Watt is K times DVF over VF times IH.

You may be wondering about that "brief delay" TMD after the heating pulse and before the final VF reading. Ideally TMD would be zero so that the junction would have no time to cool before the final "hot" VF measurement. There are practical limitations however, some related to the diode and some to the test equipment. Published specifications normally specify the TMD to be used. This is typically in tens of microseconds for small fast diodes to hundreds of microseconds for large slow diodes.

The tests described above produce a single reading for a set of test conditions that you have chosen. It is also very useful to make plots of Thermal Response vs TMD and Thermal Response vs TH. These plots, called "Cooling Plot" and "Heating Plot" respectively, are described fully in another paper by this writer.

Frothingham Electronics Corp. produces a number of testers that measure DVF and THETA as described. Software included can produce the plot mentioned. Please contact this writer if you have any questions.

What is the difference between DVF as measured on a VF40 and on an FEC200?

In an ideal (but unrealistic) DVF test, the diode and the test equipment would both be infinitely fast. So at the end of the heating pulse, the current would go instantly from IH to IM. The heating power would drop to near zero and the junction would begin to cool.

After a delay (TMD) the VF would be measured and compared with the cold reading to get DVF. The junction would of course be a bit cooler than at the end of the IH pulse because of the cooling during TMD.

In a real DVF test however, both the diode and the test equipment take a finite (and variable) time to switch from IH to IM.

The test equipment is limited by the turn-off time of the IH power supply, which, in turn is increased by any inductance in the circuit between the power supply output and the diode under test.

The diode under test may also have a substantial effect on the turn-off time. This is especially true for "standard recovery" rectifiers. One might think that a 2m S Trr would be quite small compared with a typical TMD. However, the 2m S Trr is measured under far different conditions. The Trr in a DVF test situation is **MUCH** longer. Note that the current being switched from a high Forward Current to a low Forward Current. In this situation the Trr is controlled substantially by the carrier recombination time.

We set the value of TMD accurately, measured from the time we switch off IH to the time we sample the "hot" VF at IM.

The actual effective TMD however is reduced by the switching time of the combined IH supply and the diode under test and circuit inductance.

IMPORTANT NOTE: During the switching time of IH to IM the measured VF is very unpredictable and is substantially meaningless. It is very important to avoid using a TMD in this area.

A very instructive demonstration of this effect is to use the "Cooling Plot" function of the VFS2 software and a VF40 tester. Run this plot on two rectifiers of the same or very similar size. One rectifier should be a "standard recovery" type and the other a "fast recovery". Note the significant difference between the two plots at the low TMD end. Also note what would happen if you made a single measurement of DVF on a slow rectifier at a low TMD that worked well for a fast diode.

In addition to all of this, there is an inductive coupling of the force and sense leads in long cables that effects the stability of the VF sensing circuits to acquire an accurate reading. The cables that we supply with the FEC200 are designed to minimize this effect when connected to our manual test station or to a handler where there is minimal wiring from the test clips to the ends of our cables. At some combination of IH, IM, and TMD the fixturing can be fairly critical.

THERMAL RESPONSE PLOTS

Many of the Thermal Response testers sold by FEC include software which (among other things) produces two types of plots that we call "Cooling Plot" and "Heating Plot." This paper describes these two plots and their intended applications.

COOLING PLOTS

This plots THERMAL RESPONSE vs reading delay TMD. The TMD normally used for a single reading is typically in the 50 to 100 microsecond range. The plots however start at a short time such as 10m S and extend out to tens of milliseconds. The high and low ends of the plots have different intended uses.

The low end of the plot allows you to evaluate the performance of the tester, the test fixture and device under test to find the lowest safe value of TMD to use for your screening tests.

It is desirable to use as low a TMD as possible since the junction is cooling significantly during this delay. We would use zero delay if we could but practical considerations related to the tester and to the device under test make this impossible.

When you look at a Cooling Plot you will see the Thermal Response increasing smoothly as the TMD decreases until at some point there is a marked discontinuity in the curve. You can easily see the lowest TMD that is on the smooth predictable part of the curve. Of course if you use a lower TMD, your test results will be unreliable.

The high end of the plot has a use similar to the Heating Plot so much of that will be covered in the section on Heating Plots. It is useful to note here though, that the Cooling Plot is somewhat safer to the device under test as well as faster to produce. This is so because it uses a fixed heating pulse width thereby reducing the possibility of overheating the part. The cooling time between points on the plot can be relatively short and constant over the plot. This reduces the total plotting time.

You will note when you compare the two plots that they are very nearly mirror images of one another.

The above might lead you to think there is no need for heating plots, but this is not so. Heating plots have much better resolution at the high end of the time range. When the TMD for a cooling plot is very long the junction has cooled to such a degree that DVF is very small and hard to measure.

Heating plots cause the junction to get very hot at the high end of the time scale. This results in a large easy to read DVF. If DVF gets too large, we will suspend the plot and mark the remaining programmed points as "Invalid". This protects both the device under test and the tester.

PLOT COOLING FIRST

Why? When you do a heating plot you must choose a value for TMD. The cooling plot allows you to find a low but safe value so that your heating plot will give reliable results.

WHAT ARE WE LOOKING FOR ANYWAY?

The most important use of the heating plot or the long end of the cooling plot is to find the most effective pulse width and test limit for detecting bad bonds.

Many registered specifications set the test conditions and limits that you must pass and of course you must do so, **but** passing the registered specification is no guarantee that the parts have good bonds.

For a given part number and a given manufacturing process, the difference between a good bond and an unreliable bond may be fairly small. This is particularly true if the pulse width is not optimum. Each manufacturer's process yields a unique profile of junction temperature vs heating time. Usually all of the processes result in good parts providing that the process is in good control.

An industry wide fixed Thermal Response test must necessarily be written to accept the product of most quality manufacturers. This type of conservative specification will usually fail really bad parts but will probably not be optimum for your particular process.

We suggest that you make heating plots for a reasonably large sample of your parts and examine the plots looking for a few parts whose curves show a noticeable increase in slope at some moderate pulse width. This small deviation usually occurs at about the time that the heat flow from the junction first reaches the solder bond.

We at FEC have had the opportunity over the years to work with several manufacturers who suspected that they had bonding problems. Here is what usually happens.

First we make a lot of heating plots and look for parts with a noticeable difference in the rate of increase of temperature. Some of the good and some of the suspected bad parts are cut open to confirm the correlation between the curves and the bonds.

We then setup a test with optimum pulse width and test limits to reject the bad parts.

That is good but the next part is even better. Now that we have a reliable method of evaluating the quality of the bond, the process can be adjusted to give better results. This usually results in parts that are typically better than even the formerly "good" parts.

Please contact Frothingham Electronics Corp. if you have any questions or need assistance with your Thermal Response testing.

www.frothingham.com